Towards the Design of an End-to-End Automated System for Image and Video-based Recognition

Rama Chellappa[1], Jun-Cheng Chen[3], Rajeev Ranjan[1], Swami Sankaranarayanan[1], Amit Kumar[1], Vishal M. Patel[2] and Carlos D. Castillo[4]

arXiv:1601.07883v1 [cs.CV] 28 Jan 2016

***Abstract*—Over many decades, researchers working in object recognition have longed for an end-to-end automated system that will simply accept 2D or 3D image or videos as inputs and output the labels of objects in the input data. Computer vision methods that use representations derived based on geometric, radiometric and neural considerations and statistical and structural matchers and artificial neural network-based methods where a multi-layer network learns the mapping from inputs to class labels have provided competing approaches for image recognition problems. Over the last four years, methods based on Deep Convolutional Neural Networks (DCNNs) have shown impressive performance improvements on object detection/recognition challenge problems. This has been made possible due to the availability of large annotated data, a better understanding of the non-linear mapping between image and class labels as well as the affordability of GPUs. In this paper, we present a brief history of developments in computer vision and artificial neural networks over the last forty years for the problem of image-based recognition. We then present the design details of a deep learning system for end-to-end unconstrained face verification/recognition. Some open issues regarding DCNNs for object recognition problems are then discussed. We caution the readers that the views expressed in this paper are from the authors and authors only!**

I. Introduction

Over many decades, researchers working in object recognition have longed for an end-to-end automated system that will simply accept 2D or 3D image or videos as inputs and output the labels of objects in the input data. In object recognition systems developed thus far, representations such as templates, interest points, curves, surfaces, appearance models, parts, histogram of gradients, scale invariant feature transform, stochastic models, dynamic textures and many others have been used. For recognition, statistical, syntactic recognition methods, SVMs, graph matchers, interpretation trees, and many others have been employed. Despite significant progress, the performance of these systems has not been adequate for deployment. In a parallel universe, systems based on artificial neural networks have shown much promise since the mid-eighties. Given the availability of millions of annotated data, GPUs and a better understanding of the non-linearities, DCNNs are providing much better performance on tasks such as object/face detection, object recognition, face verification/recognition than traditional computer vision methods.

In this paper, we briefly trace the history of development of object recognition methods in images and videos using methods based on traditional computer vision and artificial neural networks and then present a case study in designing an end-to-end face verification/recognition system using deep learning networks. While the performance of methods based on DCNNs is impressive, we feel that current day DCNNs are more like the Model-T cars. We conclude the paper with a summary of remaining challenges to be addressed so that deep learning networks can morph into high performance robust recognition systems.

The organization of this paper is as follows: Section II presents a brief history of developments in computer vision and artificial neural networks over the last forty years. Section III presents the design details of a deep learning system for end-to-end unconstrained face verification/recognition under the support of IARPA JANUS program. Some open issues regarding DCNNs for object recognition problems are discussed in Section IV.

II. A Brief History of Developments in Image Recognition Using Computer Vision and Neural Network-Based Methods

Since the early sixties, when Robert's edge operator was introduced, computer vision researchers have been working on designing various object recognition systems. The goal has been to design an end-to-end automated system that will simply accept 2D, 3D or video inputs and spit out the class labels or identities of objects. Beginning with template matching approaches in the sixties and seventies, methods based on global and local shape descriptors were developed. In the seventies, methods based on representations such as Fourier descriptors, moments, Markov models, and statistical pattern recognizers were developed. Even in the early years, the need for making the global recognition approaches be invariant to various transformations such as scale, rotation, etc. were recognized and appropriate transformations were suggested in the feature extraction stage; an alternative approach was to search over these transformations during the classification step. Unlike these global descriptors, local descriptors based on primitives such as line segments, arcs etc. were used in either structural or syntactic pattern recognition engines. For example, generative grammars of various types were designed

[1]Department of ECE and the Center for Automation Research, UMIACS, University of Maryland, College Park, MD USA. email: {rama,rranjan1,swamiviv,akumar14,carlos}@umiacs.umd.edu

[2]Department of Electrical and Computer Engineering at Rutgers University, Piscataway, NJ USA. email: vishal.m.patel@rutgers.edu.

[3] Department of Computer Science and the Center for Automation Research, UMIACS, University of Maryland, College Park, MD, USA. email: pullpull@cs.umd.edu.

[4] Center for automation Research, UMIACS, University of Maryland, College Park, MD, USA.

to parse the given object contour into one of many classes. More information on these developments can be found in [12], [13].

In the eighties, statistical pattern recognition methods were seen as not being able to handle occlusions or geometric representations. Graph matching or relaxation approaches became popular for addressing problems such as partial object matching. In the mid-eighties, 3D range data of objects became available leading to surface-based descriptors, jump edges (edges separating the object and background) and crease edges (edges between surfaces). These representations naturally led to graph-based or structural matching algorithms. Another approach based on interpretation trees yielded a class of algorithms for object recognition. The theory of invariants became popular with the goal of recognizing objects over large view points. More information on these developments can be found in [11], [17], [20], [24], [37], [43].

While these approaches were being developed, methods based on artificial neural networks (ANNs) made a comeback in the mid-eighties. The emergence of ANNs was largely motivated by the excitement generated by Hopfield network's ability to address the traveling salesman problem and the rediscovery of back-propagation algorithm for training the ANNs. The first international conference on neural networks held in 1987 in San Diego attracted lot more attendees than expected. Firm believers in ANNs thought that a new era had begun, and some even claimed that the demise of artificial intelligence was inevitable! The ANNs were not broadly welcomed by computer vision researchers. The reason for the ambivalence of computer vision researchers to ANNs was understandable. Computer vision researchers were brought up with the notion that representations derived from geometric, photometric as well as human vision points of view were critical for the success of object recognition systems [32]. The approach of simply feeding images into a 3-layer ANN and getting the labels out using training data was not appealing to most computer vision researchers. For one thing, it was not clear how general invariances could be integrated into ANNs, despite early attempts of Fukushima in designing the Neocognitron. ANN experiments that demonstrated object recovery from partial information were not convincing. ANN researchers relied on the theory that the 3-layer networks can represent the mapping between inputs and the class labels, given the proper amount of neurons and training data. This claim did not satisfy computer vision researchers as it was not clear how the non-linear mappings in ANNs explicitly span the variations due to pose, illumination, occlusion etc. Also, computer vision researchers were more interested in 3D object recognition problems and were not into OCRs where the ANNs were being applied. More information on these developments can be found in [4], [14], [38], [48], [51].

While the ANNs were becoming popular, a class of networks known as Convolutional Neural Networks (CNNs) was being developed by LeCun and associates. The CNNs showed much promise in the domain of OCR. The CNNs represented the idea that one can learn the representations from data using learning algorithms. The tension between learning representations directly from data vs handcrafting the representations and applying appropriate preprocessing steps has been ever present. The emergence of representations such as the Scale-invariant feature transform (SIFT), which showed an order of magnitude improvement when compared to interest points developed more than three decades ago, the histogram of gradients (HoG) operator and the local binary pattern [1] are good examples of hand-crafted features. CNNs left the feature extraction work to a learning algorithm. Irrespective of whether hand crafted or data-driven features extracted, there was a common agreement on the effectiveness of support vector machines as classifiers. More information on these developments can be found in [10], [27]–[29], [46].

Since the mid-nineties, computer vision researchers got interested in problems such as video-based tracking, surveillance and activity recognition, enabled by the availability of large volumes of video data from stationary and moving cameras. Other problems such as video stabilization, 3D modeling from multiple cameras, gait analysis etc. also were pursued. Another interesting development is the willingness of researchers to be challenged with large data sets and performance expectations. While these challenge problems were first introduced in the OCR community, they soon made their way into object detection and recognition (PASCAL VOC, ImageNet), face recognition (FRGC, LFW, IJB-A, MegaFace) and activity recognition research communities. It has become commonplace to have one or more new challenge problems to be introduced every year. Many research programs funded by the Government also introduced data sets and evaluation protocols appropriate for measuring progress in their programs. More detains on these developments can be found [5], [8], [9], [21], [22], [33], [34], [41], [45], [50].

The undaunted stalwarts of ANNs continued their quest for improving the performance by increasing the number of layers. Since the effectiveness of backpropagation algorithm was diminishing as the number of layers increased, unsupervised methods based on Boltzman machines [19] and autoencoders [2] were suggested for obtaining good initial values of network parameters which were then fed into deep ANNs.

As these developments were being made, the "Eureka" moment came about when DCNNs were first deployed for the ImageNet challenge a mere four years back. The performance improvements obtained by DCNNs [26] for the ImageNet challenge were quite good. The power of depth, the availability of GPUs and large annotated data, replacement of traditional sigmoidal nonlinearities by Rectified Linear Units (ReLU) and drop out strategies were embodied in the network now known as AlexNet [26]. The life of computer vision researchers has not been the same since!

The success of AlexNet motivated researchers from companies and numerous universities [39], [42], [44] to design various versions of DCNNs by changing the number of layers, the amount of training data being used and modifications to the nonlinearities, etc. The tables shown below document the improvements on object detection, image classification and face verification over the last several years.

While some may be dismayed by the reemergence of the so called "black box" approach to computer vision problems such as object detection/recognition and face verifica-

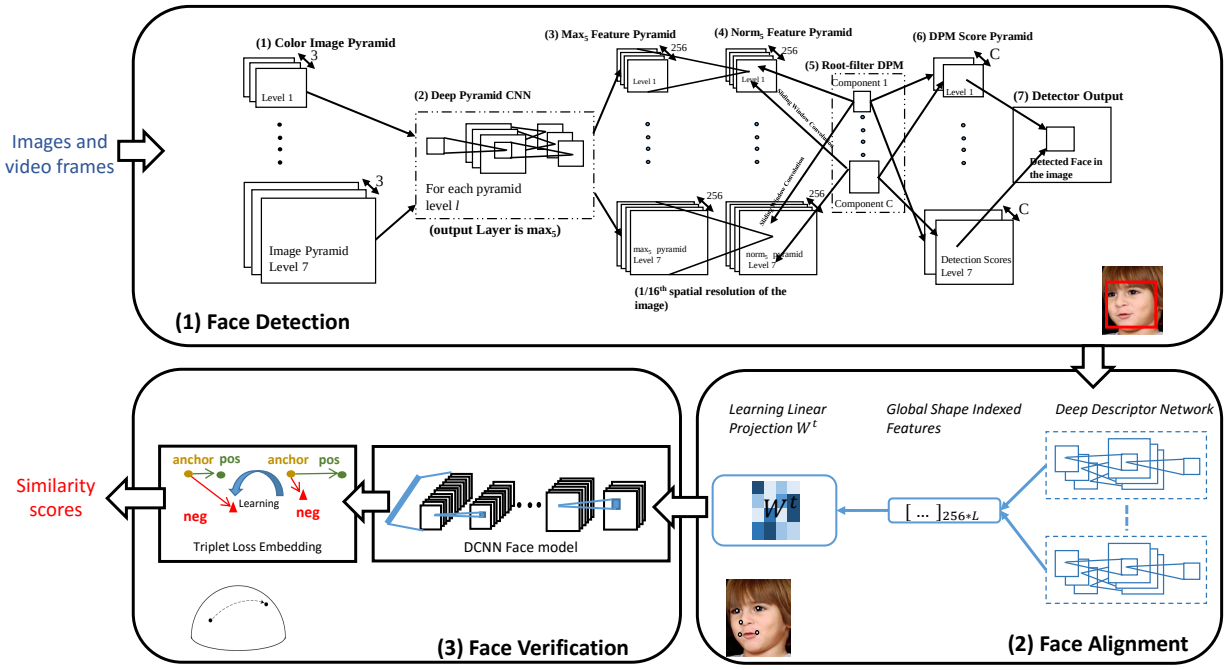

Fig. 1. An overview of the proposed end-to-end DCNN-based face verification system [7].

Rank	Name	Error rate	Description
1	Google	0.06656	Deep learning
2	Oxford	0.07325	Deep learning
3	MSRA	0.08062	Deep learning

TABLE I
TOP IMAGENET 2014 IMAGE CLASSIFICATION CHALLENGE RESULTS (HTTP://IMAGE-NET.ORG/CHALLENGES/LSVRC/2014/).

Rank	Name	Error rate	Description
1	Google	0.43933	Deep learning
2	CHUK	0.40656	Deep learning
3	DeepInsight	0.40452	Deep learning
4	UvA-Euvision	0.35421	Deep learning
5	Berkeley Vision	0.34521	Deep learning

TABLE II
TOP IMAGENET 2014 OBJECT DETECTION CHALLENGE RESULTS (HTTP://IMAGE-NET.ORG/CHALLENGES/LSVRC/2014/).

Method	Result
DeepID2	0.9915 ± 0.0013
TCIT	0.9333 ± 0.0124
DeepID2+	0.9947 ± 0.0012
betaface.com	0.9808 ± 0.0016
DeepID3	0.9953 ± 0.0010
insky.so	0.9551 ± 0.0013
Uni-Ubi	0.9900 ± 0.0032
FaceNet	0.9963 ± 0.0009
Tencent-BestImage	0.9965 ± 0.0025
Baidu	0.9977 ± 0.0006
AuthenMetric	0.9977 ± 0.0009
MMDFR	0.9902 ± 0.0019
CW-DNA-1	0.9950 ± 0.0022
Sighthound	0.9979 ± 0.0003

TABLE III
TOP MEAN CLASSIFICATION ACCURACY AND STANDARD ERROR OF THE MEAN ON THE LFW DATASET (HTTP://VIS-WWW.CS.UMASS.EDU/LFW/RESULTS.HTML).

tion/recognition, the simple fact of life is that it is hard to argue against performance. What is comforting though, the original issues in object recognition problems, such as robustness to pose, illumination variations, degradations due to low-resolution, blur and occlusion still remain. While using hundreds of million face images of millions subjects has produced some of the best results in LFW face challenge, the recently reported recognition performance of a DCNN with 500 million faces of 10 million subjects in the low seventies for the MegaFace challenge, clearly argues for much more work to be done.

Over three decades, the author's group explored both traditional computer vision approaches and ANN-based approaches for a variety of problems in computer vision including object/face detection, face verification/recognition. An earlier attempt [30] on designing FLIR ATRs using CNNs undertaken by the author's group yielded discouraging results as the number of layers used was insufficient. Although sufficient training data was available for a ten-class ATR problem, not having GPUs hampered the development of deeper CNNs. Over the last eighteen months, under the support of IARPA JANUS program for unconstrained face verification/recognition, the authors have designed an end-to-end automatic face verification and recognition system using DCNNs for face detection [35], face alignment [7] and verification/recognition stages [6], [7]. This has been made possible due to the availability of affordable GPU cards, large annotated face data sets and implementations of CNNs such as Caffe [23], Torch, CUDa-

ConvNet, Theano and TensorFlow. Even high-school students are able to build vision applications using DCNNs [36]! While this is an exciting development for most researchers interested in designing end-to-end object recognition system, we feel we are at a stage similar to when Model-T was introduced many decades ago.

III. UMD Face Verification System

In this section, we give a brief overview of the UMD face verification system [6], [7] based on DCNNs. Figure 1 shows the block diagram of the overall verification system. First, faces are detected in each image and video frame. Then, the detected faces are aligned into canonical coordinates using the detected landmarks. Finally, face verification is performed using the low-dimensional features obtained using a triplet loss formulation to compute the similarity between a pair of images or videos. In what follows, we describe each of these blocks in detail.

A. Face Detection

Faces in the images/videos are detected using the Deep Pyramid Deformable Parts Model for Face Detection (DP2MFD) approach presented in [35]. This method consists of two main stages. The first stage generates a seven level normalized deep feature pyramid for any input image of arbitrary size. The second stage is a linear SVM, which takes these features as input to classify each location as face or non-face, based on their scores. It was shown in [35] that this detector is robust to not only pose and illumination variations but also to different scales. Furthermore, DP2MFD was shown to perform better than an earlier version of Region with CNN (RCNN) based detector [15]. Figure 2 shows some detected faces by DP2MFD.

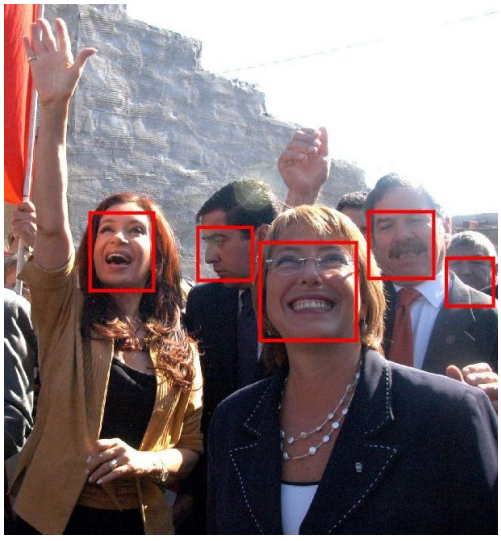

Fig. 2. Sample detection results on an IJB-A image using the deep pyramid method.

B. Landmark Localization

In the second step of our verification system, facial landmarks are detected to align the detected faces. The landmark detection problem can be viewed as a regression problem, where beginning with the initial mean shape, the target shape is reached through regression. First, deep features are extracted of a patch around a point of the shape. Then, the cascade regression in which the output generated by the first stage is used as an input for the next stage is used to learn the regression function. After the facial landmark detection is completed, each face is aligned into the canonical coordinate with similarity transform using the 3 landmark points (i.e. the center of each eye and the base of the nose). After alignment, the face image resolution is 100×100 pixels. Examples of detected landmarks results are shown in Figure 3.

Fig. 3. Sample facial landmark detection results [7].

C. Face Verification

Once the faces are detected and aligned, we train a DCNN network using the architecture presented in [6]. The network is trained using the CASIA-WebFace dataset [49]. The dimensionality of the input layer is $100 \times 100 \times 3$ for gray-scale images. The network includes 10 convolutional layers, 5 pooling layers, and 1 fully connected layer. The final dimensionality of the feature is 320. More details of the DCNN network can be found in [6].

To further improve the performance of our deep features, we obtain a low-dimensional discriminative projection of the deep features that is learnt using the training data provided for each split of IJB-A (section D below describes the IJB-A data set). The output of the procedure is an embedding matrix $\mathbf{W} \in \mathbf{R}^{M \times n}$ where M is the dimensionality of the deep descriptor, which is 320 in our case and we set $n = 128$, thus achieving dimensionality reduction in addition to an improvement in performance.

The objective of this method is to push similar pairs together and dissimilar pairs apart in the low-dimensional space. For learning $\mathbf{W}$, we solve an optimization problem based on constraints involving triplets - each containing two similar samples and one dissimilar sample. The triplets are generated using the label information provided with the training data for each split. This optimization is solved using stochastic gradient descent and the entire procedure takes 3-5 minutes per split using a standard C++ implementation.

D. Experimental Results

To show the performance of the UMD DCNN-based face verification system, we highlight some verification results on the challenging IARPA JANUS Benchmark A (IJB-A) [25] and its extended version JANUS Challenging set 2 (JANUS CS2) dataset [7]. The JANUS CS2 dataset contains not only the sampled frames and images in the IJB-A, but also the original videos. In addition, the JANUS CS2 dataset[1] includes considerably more test data for identification and verification problems in the defined protocols than the IJB-A dataset. In both the IJB-A and the CS2 data set the verification and recognition protocols are set around templates, which are groups of one or more images of the same person. Verification is template to template, recognition is template to gallery (of templates).

Table IV summarizes the scores produced by different face verification methods on the IJB-A dataset. Among the methods compared in Table IV include, a DCNN-based method [47], Fisher vector-based method [40], and one commercial off-the-shelf matcher, COTS [25] which are tested in a fully automatic setup. The DCNN that performs verification uses data processed by the automated face preprocessing components described in previous subsections. This is followed by fine-tuning, triplet embedding, and testing steps.

Since the system works end-to-end, we have devised two ways of handling the situation where we are unable to detect any of the faces in the images of given template:

- **Setup 1:** Under this setup, our verification and recognition accuracy is measured only over images we can process. The philosophy for this setup is: assume average performance for any images we cannot process. For verification: if in a template to template comparison in which we cannot detect a face in any of its images, we ignore this template to template comparison. For recognition: if we cannot process any of the images of an individual to compare versus a gallery we do not include this probe in the computation of recognition rates.
- **Setup 2:** Under this setup, our verification and recognition accuracy is measured in a pessimistic fashion: we are forced to make a decision even if we have not been able to detect faces. In this case, for verification we set the score to the lowest possible similarity (when we're unable to process any of the images in one of the two templates being compared). For recognition experiments, we set the rank as the highest possible rank, when we're unable to process all the images in the probe template.
- **Setup 3:** In this setup, we include our earlier results [7] for the purpose of comparison.

The following lists the differences of the components used in the current work from **Setup 3**:

- We adopt the same network architecture presented in [7]; However, for the training data, we use RGB color and larger face region (*i.e.*, crop 125 $\times$ 125-pixel face regions and resize them to 100 $\times$ 100 pixels.) during alignment instead of gray scale and smaller face regions.
- We replace the joint Bayesian metric with triplet loss embedding.
- Complete implementation in C++ instead of MATLAB.
- **Setup 3** has one extra component, face association, to detect and track faces across frames. The association component requires one of ground truth face bounding boxes as the initialization to specify which face track to be used for comparison. Thus, if all the faces are not detected within a template, it still has one face for comparison as compared to other setups. (*i.e.* However, the faces are difficult ones, and most of them are in extreme pose and illumination variations. Thus, including these faces does not improve the overall performance a lot.)

Figure 4 shows the ROC curves and the CMC curves corresponding to different methods on the JANUS CS2 dataset, respectively for verification and identification protocols. The corresponding scores are summarized in Table V. From the ROC and CMC curves, we see that the DCNN method performs better than other competitive methods. This can be attributed to the fact that the DCNN model does capture face variations over a large dataset and generalizes well to a new small dataset. In addition, the performance of the proposed automatic system degrades only slightly as compared to the one using the manual annotations. This demonstrates the robustness of each component of our system.

IV. Open Issues

Given sufficient number of annotated data and GPUs, DCNNs have been shown to yield impressive performance improvements. Still many issues remain to be addressed to make the DCNN-based recognition systems robust and practical. These are briefly discussed below.

- **Reliance on large training data sets:** As discussed before, one of the top performing networks in the MegaFace challenge needs 500 million faces of about 10 million subjects. Such large annotated training set may not be always available (e.g. expression recognition, age estimation). So networks that can perform well with reasonable-sized training data are needed.
- **Invariance:** While limited invariance to translation is possible with existing DCNNs, networks that can incorporate more general 3D invariances are needed.
- **Training time:** The training time even when GPUs are used can be several tens to hundreds of hours, depending on the number of layers used and the training data size. More efficient implementations of learning algorithms, preferably implemented using CPUs are desired.
- **Number of parameters:** The number of parameters can be several tens of millions. Novel strategies that reduce the number of parameters need to be developed.
- **Handling degradations in training data:** : DCNNs robust to low-resolution, blur, illumination and pose variations, occlusion, erroneous annotation, etc. are needed to handle degradations in data.
- **Domain adaptation of DCNNs:** While having large volumes of data may help with processing test data from

[1]The JANUS CS2 dataset is not publicly available yet.

IJB-A-Verif	[47]	DCNN (setup 1)	DCNN (setup 2)	DCNN (setup 3)
FAR=1e-2	0.732 ± 0.033	0.8312 ± 0.0350	0.7810 ± 0.0316	0.776 ± 0.033
FAR=1e-1	0.895 ± 0.013	0.9634 ± 0.0049	0.9006 ± 0.0077	0.936 ± 0.01
IJB-A-Ident	[47]	DCNN (setup 1)	DCNN (setup 2)	DCNN (setup 3)
Rank-1	0.820 ± 0.024	0.8990 ± 0.0105	0.8378 ± 0.0142	0.834 ± 0.017
Rank-5	0.929 ± 0.013	0.9706 ± 0.0075	0.9073 ± 0.0119	0.922 ± 0.011
Rank-10	N/A	0.9821 ± 0.0053	0.9219 ± 0.0094	0.947 ± 0.011

TABLE IV
RESULTS ON THE IJB-A DATASET. THE TAR OF ALL THE APPROACHES AT FAR=0.1 AND 0.01 FOR THE ROC CURVES. THE RANK-1, RANK-5, AND RANK-10 RETRIEVAL ACCURACIES OF THE CMC CURVES. WE REPORT AVERAGE AND STANDARD DEVIATION OF THE 10 SPLITS.

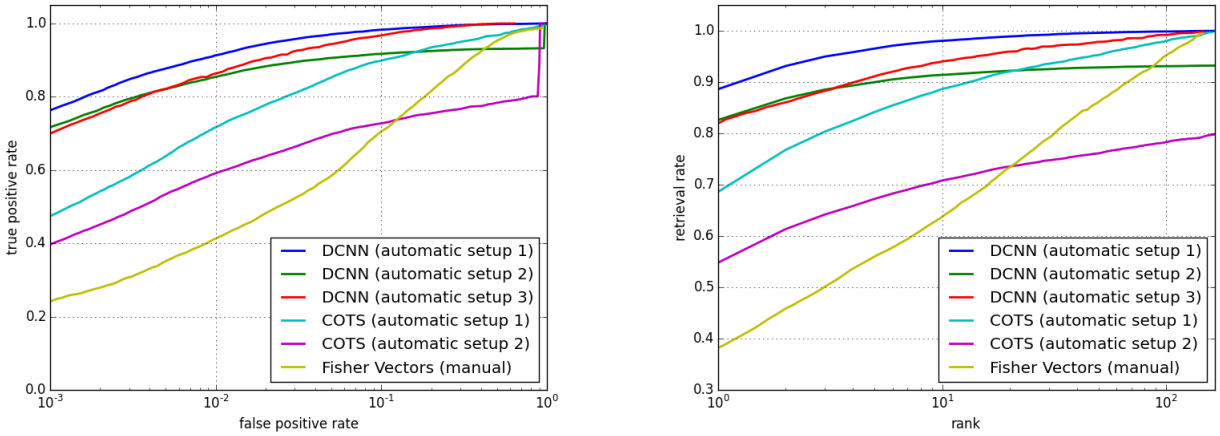

(a) (b)

Fig. 4. Results on the JANUS CS2 dataset. (a) the average ROC curves and (b) the average CMC curves. Our curves are the average curve of the 10 splits.

JANUS-CS2-Verif	COTS (setup 1)	COTS (setup 2)	FV [40]	DCNN (setup 1)	DCNN (setup 2)	DCNN (setup 3)
FAR=1e-2	0.7167 ± 0.0168	0.5917 ± 0.0154	0.411 ± 0.081	0.9132 ± 0.0059	0.8538 ± 0.008	0.861 ± 0.014
FAR=1e-1	0.8986 ± 0.0099	0.7272 ± 0.0133	0.704 ± 0.028	0.9829 ± 0.0038	0.9170 ± 0.0074	0.963 ± 0.007
JANUS-CS2-Ident	COTS (setup 1)	COTS (setup 2)	FV [40]	DCNN (setup 1)	DCNN (setup 2)	DCNN (setup 3)
Rank-1	0.6861 ± 0.0156	0.5481 ± 0.0150	0.381 ± 0.018	0.8862 ± 0.0089	0.8262 ± 0.0100	0.82 ± 0.014
Rank-5	0.8417 ± 0.0111	0.6724 ± 0.0141	0.559 ± 0.021	0.9653 ± 0.0057	0.9000 ± 0.0086	0.91 ± 0.01
Rank-10	0.8865 ± 0.0091	0.7081 ± 0.0125	0.637 ± 0.025	0.9803 ± 0.0047	0.9139 ± 0.0081	0.938 ± 0.01

TABLE V
RESULTS ON THE JANUS CS2 DATASET. THE TAR OF ALL THE APPROACHES AT FAR=0.1 AND 0.01 FOR THE ROC CURVES. THE RANK-1, RANK-5, AND RANK-10 RETRIEVAL ACCURACIES OF THE CMC CURVES. WE REPORT AVERAGE AND STANDARD DEVIATION OF THE 10 SPLITS.

a different distribution than that of the training data, systematic methods for adapting the deep features to test data are needed.

- **Theoretical considerations:** While DCNNs have been around for a few years, detailed theoretical understanding is just starting to develop [3], [16], [18], [31]. Methods for deciding the number of layers, neighborhoods over which max pooling operations are performed are needed.
- **Incorporating domain knowledge:** The current practice is to rely on fine tuning. For example, for the age estimation problem, one can start with one of the standard networks such as the AlexNet and fine tune it using aging data. While this may be reasonable for somewhat related problems (face recognition and facial expression recognition), such fine tuning strategies may not always be effective. Methods that can incorporate context may make the DCNNs more applicable to a wider variety of problems.
- **Memory:** The DCNNs in use currently are memoryless, in that they cannot predict based on the data they have seen before. This reduces their ability to process videos. Although Recurrent CNNs are on the rise, they still consume a lot of time and memory for training and deployment. Efficient DCNN algorithms are needed to handle videos and other data streams as blocks.

ACKNOWLEDGMENTS

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2014-14071600012. The views and conclusions contained herein are those of the authors and should not

be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

REFERENCES

[1] T. Ahonen, A. Hadid, and M. Pietikainen. Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):2037–2041, 2006.

[2] Y. Bengio. Learning deep architectures for ai. *Found. Trends Mach. Learn.*, 2(1):1–127, Jan. 2009.

[3] J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013.

[4] G. Carpenter and S. Grossberg. *Pattern Recognition by Self-Organizing Neural Networks*. A Bradford Book, 1991.

[5] R. Chellappa, A. Sankaranarayanan, A. Veerraghavan, and P. Turaga. *Statistical Models and Methods for Video-based Tracking and Recognition*. Now Publishers Inc, 2010.

[6] J.-C. Chen, V. M. Patel, and R. Chellappa. Unconstrained face verification using deep cnn features. In *IEEE Winter conference on Applications of Computer Vision*, 2016.

[7] J.-C. Chen, R. Ranjan, A. Kumar, C.-H. Chen, V. M. Patel, and R. Chellappa. An end- to-end system for unconstrained face verification with deep convolutional neural networks. In *IEEE International Conference on Computer Vision (ICCV) workshop on ChaLearn Looking at People (ChaLearn LaP)*, 2016.

[8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009.

[9] M. Everingham, L. V. Gool, C. K. I.Williams, J.Winn, and A. Zisserman. The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.

[10] C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Learning hierarchical features for scene labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1915–1929, Aug 2013.

[11] O. Faugeras. *Three-dimensional Computer Vision: A Geometric View Point*. The MIT Press, 1993.

[12] K. S. Fu. *Syntactic Pattern Recognition and Applications*. Prentice-Hall, 1981.

[13] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, 1990.

[14] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4):93–202, 1980.

[15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 580–587, 2014.

[16] R. Giryes, G. Sapiro, and A. M. Bronstein. On the stability of deep networks. *arXiv preprint arXiv:1412.5896*, 2014.

[17] W. Grimson. *Object Recognition by Computer: The Role of Geometric Constraints*. MIT Press, Cambridge, MA, 1990.

[18] B. D. Haeffele and R. Vidal. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.

[19] G. E. Hinton and T. J. Sejnowski. *Learning and relearning in Boltzmann machines in Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. MIT Press, Cambridge, MA, 1986.

[20] B. K. P. Horn. *Robot Vision*. The MIT Press, Cambridge, MA, 1986.

[21] K. Ikeuchi. *Computer Vision: A Reference Guide*. Springer, 2014.

[22] M. Isard and A. Blake. Condensation-conditional density propagation for visual tracking. *International Journal on Computer Vision*, 29(1):5–28, 1998.

[23] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM International Conference on Multimedia*, pages 675–678, 2014.

[24] T. Kanade. *Three Dimensional Vision*. Kluwer Academic Publishers, Boston, MA, 1987.

[25] B. F. Klare, B. Klein, E. Taborsky, A. Blanton, J. Cheney, K. Allen, P. Grother, A. Mah, M. Burge, and A. K. Jain. Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[26] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

[27] Y. LeCun. *Mod'eles connexionistes de lapprentissage*. PhD thesis, Universite de Paris VI, 1987.

[28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[29] Y. LeCun, F.-J. Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 97–104, 2004.

[30] B. Li, R. Chellappa, Q. Zheng, S. Der, N. M. Nasrabadi, L. Chan, and L. Wang. Experimental evaluation of flir atr approaches - a comparative study. *Computer Vision and image understanding*, 84:5–24, 2001.

[31] S. Mallat. Understanding deep convolutional networks. *arXiv preprint arXiv:1601.04920*, 2016.

[32] D. Marr. *Vision*. The MIT Press, Cambridge, MA, 1982.

[33] M. Nixon, T. Tan, and R. Chellappa. *Human Identification Based on Gait*. Springer, 2005.

[34] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 947–954 vol. 1, June 2005.

[35] R. Ranjan, V. M. Patel, and R. Chellappa. A deep pyramid deformable part model for face detection. In *IEEE International Conference on Biometrics: Theory, Applications and Systems*, 2015.

[36] R. Ranjan, S. Zhou, J.-C. Chen, A. Kumar, A. Alavi, V. M. Patel, and R. Chellappa. Unconstrained age estimation with deep convolutional neural networks. In *IEEE International Conference on Computer Vision (ICCV) workshop on ChaLearn Looking at People (ChaLearn LaP)*, 2016.

[37] A. Rosenfeld and A. Kak. *Digital Picture Processing Vol 1 and 2*. Academic Press, 1982.

[38] D. E. Rumelhart, J. L. McClelland, and the PDP Research Group. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol 1*. Cambridge: MIT Press, 1986.

[39] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, June 2015.

[40] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher vector faces in the wild. In *British Machine Vision Conference*, volume 1, page 7, 2013.

[41] A. Smeulders, D. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: An experimental survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1442–1468, July 2014.

[42] Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2892–2900, June 2015.

[43] R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer, 2010.

[44] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, June 2014.

[45] P. Turaga, R. Chellappa, V. Subrahmanian, and O. Udrea. Machine recognition of human activities: A survey. *IEEE Transactions on Circuits and Systems for Video technology*, 18:1473–1488, Nov. 2008.

[46] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 1995.

[47] D. Wang, C. Otto, and A. K. Jain. Face search at scale: 80 million gallery. *arXiv preprint arXiv:1507.07242*, 2015.

[48] P. Werbos. *The roots of backpropagation: from ordered derivatives to neural networks and political forecasting*. Wiley-Interscience, New York, NY, 1994.

[49] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.

[50] W. Zhao, R. Chellappa, J. Phillips, and A. Rosenfeld. Face recognition in still and video images: A literature surrey. *ACM Computing Surveys*, 35:399–458, Dec. 2003.

[51] Y. Zhou and R. Chellappa. *Artificial Neural Networks for Computer Vision*. Springer-Verlag, 1991.